
Job submission overview



Open Science Grid



THE UNIVERSITY OF
CHICAGO

Marco Mambelli – marco@hep.uchicago.edu

August 9 2011

OSG Summer Workshop

TTU - Lubbock, TX

Campus and Production Grid

■ Campus Grid

- ❑ Direct trust relationship
- ❑ Local authentication (username and password)
- ❑ Accessible as a single Condor pool

■ Production Grid

- ❑ OSG and VO mediated trust relationship
- ❑ Authentication via grid certificates (x509)
- ❑ Varied resources accessed via GRAM protocol
- ❑ Different tools to submit jobs


Job classification

- HPC High Performance Computing
 - Massively parallel
 - Specialized resources
- HTC High Throughput Computing
 - Many independent jobs
 - Common denominator across resources
- HTPC High Throughput Parallel Computing
 - Take advantage of multi-core architectures
 - Multiple threads on a single node
- All jobs in OSG should have minimal assumptions about the resources


High Throughput Computing

- Opposed to High Performance Computing
- Simple serial jobs, Monte Carlo simulation, parameter sweep (pleasantly parallel)
- Complex workflows with data and/or job dependencies
- User is decoupled from the resource
- Must be able to express the work in a highly portable way
- Cannot care about exactly where pieces are executed
- May not have SSH access to resources -> many implications
- Must be able to respect the community rules of a shared infrastructure; “plays well with others”
- Minimize time between job-submit and job-complete
- and...ensembles of small node count MPI jobs (HTPC)

Local vs Distributed



Local




Distributed

- Submission, wait in queue, execution, completion
- Variety of resources
- Many possibilities for failures (more complex)
- Different order
- Big latency
- Black boxes
- Difficult end-to-end monitoring


Direct submission vs Brokering

- Competing for resources



Direct submission


- Coordinate the scheduling
- There are still others



Brokering

Single job vs Workflow



- Direct acyclic graph
- Workflow manager
 - Runs jobs in order
 - Retries



Direct execution vs Pilot jobs

■ Pilots

- ❑ Separate user job from grid job
- ❑ Have some overhead
- ❑ Can perform some initial test
- ❑ Uniform (enhanced) environment
- ❑ Delayed scheduling



Pilot system: Pilot vs Direct job submission

Pilot

- Submitted by factories
 - Central (GiWMS, autopilot)
 - Cluster factories
- Managed by factories
- Scheduled centrally (allows late scheduling)
- Code to support VO job execution
- Submitted continuously
- Partially accounted
 - no big deal if some fail



Direct job

- Submitted by users or production managers
- Managed by workflow manager (or user)
- Scheduled by local resource managers (cluster queues)
- Code for both support and VO tasks
- Submitted when needed
- Fully accounted
 - error statistics

Pilot system: Pilot vs VO job

Pilot

- Submitted by factories
 - Central (GiWMS, autopilot)
 - Cluster factories
- Managed by factories
- Code to support job execution
- Submitted continuously
- Partially accounted
 - no big deal if some fail

VO job (on pilot)

- Submitted by users or production managers
- Managed by Panda/GiWMS Server
- Runs VO software
- Submitted when needed
- Fully accounted
 - error statistics

Low level job management

■ Tools

- ❑ LRM: Condor, PBS, LSF, SGE, ...
- ❑ Globus GRAM

■ They are

- ❑ More flexible
- ❑ A lot of different tools
- ❑ Steep learning curve
- ❑ Technology may change (CREAM)

More advanced tools

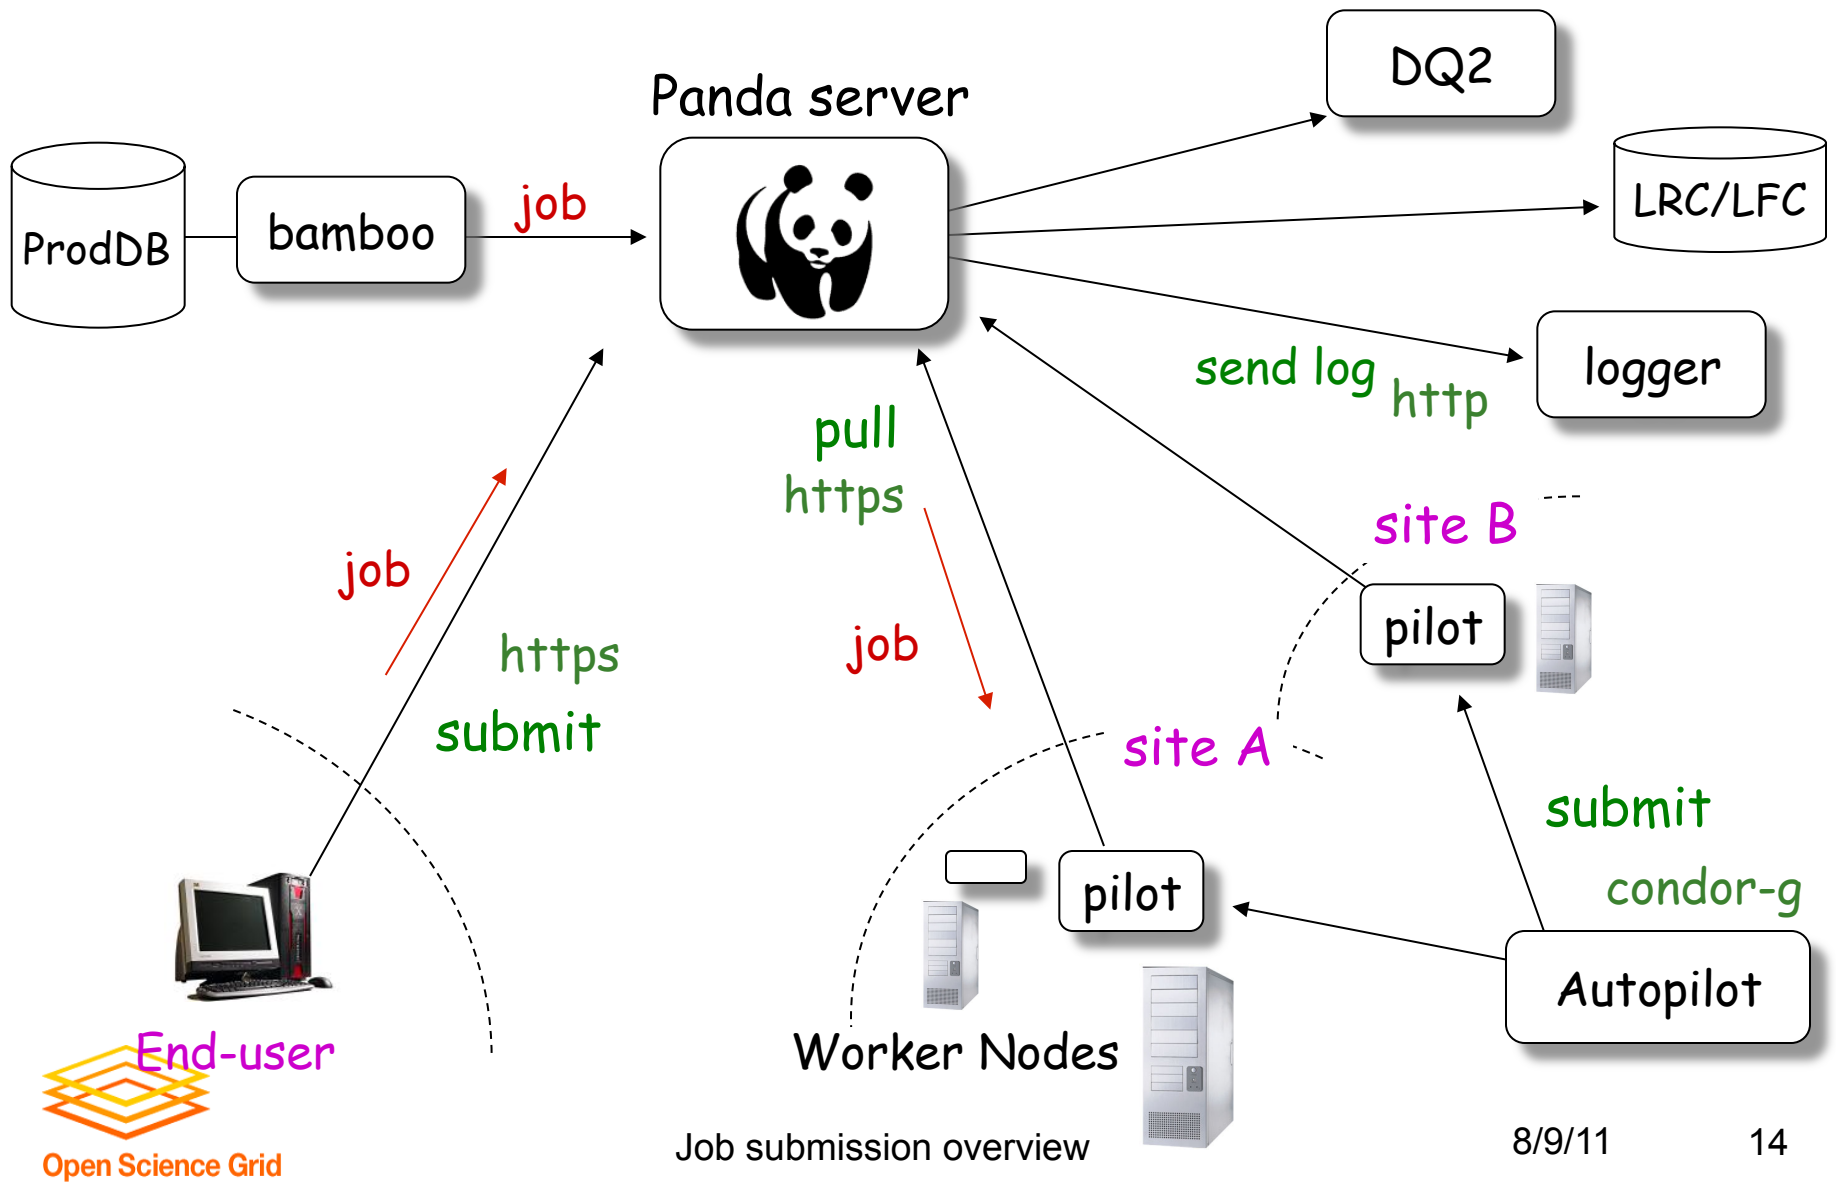
- Condor-G
 - Uniform interface, workflow management (Dagman)
- OSG MM
 - Brokering, uniform interface (C-G), workflow mgm. (Dg)
- **Glidein WMS**
 - **Brokering, uniform interface (C-G), workflow mgm. (Dg)**
- Panda
 - Monitoring, brokering, uniform interface (C-G), workflow mgm. (Dg), integrated with VO software
- CRAB
 - integrated with VO software



PANDA

- PANDA = Production ANd Distributed Analysis system
 - Designed for analysis as well as production
 - Project started Aug 2005, prototype Sep 2005, production Dec 2005
 - Works both with OSG and EGEE middleware
- A single task queue and pilots
 - Apache-based Central Server
 - Pilots retrieve jobs from the server as soon as CPU is available → late scheduling
- Highly automated, has an integrated monitoring system
- Integrated with ATLAS Distributed Data Management (DDM) system
- Not exclusively ATLAS: CHARMM, OSG ITB

Panda System



Panda Monitor: production

Panda monitor
Times are in UTC

Wiki
Shift elog
Bugs
CERN elog

Jobs - search
Recent running,
activated, waiting,
assigned, defined,
finished, failed jobs
Select analysis, prod,
install, test jobs

Quick search

Job
Dataset
Task request
Task status
File

Summaries

Blocks: days

Errors: days

Nodes: days

Daily usage

Tasks - search

Generic Task Req

EvGen Task Req

CITBsim Task Req

Task list

New Tag

Bug Report

Datasets - search

Dataset browser

Aborted MC datasets

Panda subscriptions

Datasets Distribution

DDM Req

Req list

AODs

Panda Production Operations Dashboard

Panda shift guide [calendar](#) [mailing list](#)
ADCoS [twiki](#) [calendar](#) [mailing list](#)
Production task support [mailing list](#)

[Click for help](#)

Servers: BNL:OK BNLdev:OK CERN:OK Logger:OK Bamboo:OK


Active tasks: CA:14 CERN:1 DE:15 ES:4 FR:22 IT:12 NL:1 TW:15 UK:6 US:18

Bamboo [task brokerage](#), [job submissions](#), [status](#) over last 12 hours

Jobs updated >12 hrs ago: activated:5950 running:none

Jobs updated >36 hrs ago: transferring:12

Ganglia World Wide Summary



<http://panda.cern.ch/server/pandamon/query?dash=prod>

Space available at sites:

Site GB As of (UTC)

US

AGLT2-condor 6921 09-17 02:19

ANALY_AGLT2-condor 6933 09-17 01:10

ANALY_MWT2-pbs 66815 09-17 04:27

ANALY_NET2-pbs 41342 09-17 04:21

ANALY_SLAC-lsf 16401 09-17 04:50

ANALY_SWT2_CPB-pbs 144839 09-17 01:08

BU_ATLAS_Tier2 41342 09-17 03:39

BU_ATLAS_Tier2o 41342 09-17 04:49

HU_ATLAS_Tier2 41342 09-17 02:40

IU_OSG-pbs 76176 09-17 04:49

MWT2_IU-pbs 76176 09-17 04:41

MWT2_UC-pbs

OU_OCHEP_SWT2

SLACKPD-lsf

SWT2_CPB-pbs

UC_ATLAS_MWT2

Other SEs reporting

ALBERTA-LCG2 PF

FZK-LCG2 MCDISK

IN2P3-CC MCDISK

INFN-T1 MCDISK













NDGF-T1 MCDISK

PIC MCDISK

RAL-LCG2 MCDISK

SARA-MATRIX MCI

IT	566	8372	09-17 04:50	255	0	15	0	3417	0	1776	50	443 / 0	2003	668	25%	0%	25%
NL	2	228	09-16 23:24	688	0	227	0	0	0	0	0	0 / 0	0	1	100%	0%	100%
UK	870	8439	09-17 04:50	804	0	0	0	2140	0	1599	46	1236 / 0	2905	513	15%	0%	15%
US	1526	13514	09-17 04:50	1275	0	5	0	152	0	4113	230	2526 / 12	6263	225	3%	0%	3%
TW	468	7643	09-17 04:50	241	0	10	0	3241	1	1624	29	267 / 0	2444	27	1%	0%	1%

	US Sites	Nodes	Jobs	Latest	Pilots (3hrs)	defined	assigned	waiting	activated	sent	running	holding	transferring	finished	failed tot	trf	other
	Site Name	1526	13514	09-17 04:50	1275	0	5	0	152	0	4113	230	2526 / 12	6263	225	3%	0% 3%
	AGLT2 	114	2009	09-17 04:50	197	0	0	0	0	0	529	48	935 / 0	397	100	20%	0% 20%
	BNL ATLAS 1 	318	2580	09-17 04:50	184	0	0	0	1	0	914	46	0 / 0	1605	14	1%	0% 1%
	BNL ATLAS 2 	0	0		0	0	0	0	0	0	0	0	0 / 0	0	0		
	BU ATLAS_Tier2 	12	44	09-17 04:46	0	0	0	0	13	0	22	6	2 / 0	0	1	100%	0% 100%
	BU ATLAS_Tier2o 	85	845	09-17 04:50	85	0	0	0	0	0	299	5	88 / 0	429	24	5%	5% 0%
	GLOW-ATLAS 	47	160	09-17 04:40	0	0	0	0	73	0	1	7	76 / 0	2	1	33%	0% 33%
	HU ATLAS_Tier2 	24	74	09-17 04:44	83	0	5	0	0	0	0	0	0 / 0	0	69	100%	0% 100%
	IU_OSG 	17	232	09-17 04:50	19	0	0	0	6	0	124	0	41 / 0	60	1	2%	0% 2%
	LTU CCT 	0	0		0	0	0	0	0	0	0	0	0 / 0	0	0		
	MWT2_IU 	118	1569	09-17 04:50	283	0	0	0	0	0	608	32	166 / 0	762	1	0%	0% 0%
	MWT2_UC 	184	2916	09-17 04:50	14	0	0	0	21	0	761	20	567 / 0	1545	2	0%	0% 0%
	OU OCHEP SWT2 	60	455	09-17 04:50	102	0	0	0	0	0	161	6	89 / 0	219	0	0%	0% 0%

Panda Monitor: Dataset browser

[Configuration](#)[Production](#)[Clouds](#)[DDM](#)[PandaMover](#)[AutoPilot](#)[Sites & Grids](#)[Analysis](#)[Physics data](#)[Usage & Quotas](#)[Plots](#)[ProdDash](#)[DDMDash](#)

[Update](#)

[Panda monitor](#)**DQ2 dataset browser for csc datasets**[Click for help](#)

Dataset lists last updated 157 min ago

[Quick guide, twiki](#)

Jobs - [search](#)

Recent [running](#),
[activated](#), [waiting](#),
[assigned](#), [defined](#),
[finished](#), [failed](#) jobs

Select [analysis](#),
[production](#), [test](#) jobs

Quick search

Job

Dataset

Task

File

Summaries

Blocks: days

Errors: days

Nodes: days

[Daily usage](#)

Tasks - [search](#)

[Generic Task Req](#)

[EvGen Task Req](#)

[CTBsim Task Req](#)

[Task list](#)

[Task browser](#)

Datasets - [search](#)

[Dataset browser](#)

[New datasets](#)

[Aborted MC datasets](#)

Select a project:

Or (the old way) select a dataset category *Counts are totals, exclusive of selections*

Category	Count	Description
All	147033	All datasets
DBrelease	12	DB release datasets
M3	386	M3 cosmics run
M4	4210	M3 cosmics run
conditions	35	Datasets for conditions data files
csc	4411	Computing system commissioning production
ctb	613	Combined testbeam production
dc2	6	Data Challenge 2 production
larg	53	LAr commissioning
mc	6135	MC validation production
other	71	Everything else
rome	210	Rome physics workshop production
testpanda	6439	Panda test datasets
tile	52	Tilecal commissioning
user	58456	User datasets
validation	642	Validation samples (testIdeal* etc)

Choose a site if you want to restrict dataset listings to site-resident datasets

CANADA	CERN	FRANCE	GERMANY	ITALY	NDGF	NL	SPAIN	TAIWAN	UK
ALBERTA	CERNCAF	FRTIER2S	CSCS	CNAF	IJST2	IHEP	IFAE	ASGC	RAL
MCGILL	CERNPROD	LPC	CYF	CNAFDISK	NDGFDISK	ITEP	IFIC	ASGCDISK	Rutherford
MONTREAL	TIER0DISK	LAL	DESY-HH	CNAFTAPE	NDGFT1	JINR	IFICDISK	ASGCDISK_V2	Rutherford
SFU	TIER0TAPE	SACLAY	DESY-ZN	CNAFTTEST	NDGFT1DISK	NIKHEF	IFICTAPE	ASGCTAPE	SLAC
TORON		LPNHE	FZK	LNF	NDGFT1TAPE	PNPI	LIP-COIMBRA	ASGCTAPE_V2	UK

Panda Monitor: error reporting

Panda monitor

Times are in UTC

[Wiki](#)
[Bugs](#)

[Shift elog](#)
[CERN elog](#)

Jobs - [search](#)

Recent [running](#), [activated](#), [waiting](#), [assigned](#), [defined](#), [finished](#), [failed](#) jobs

Select [analysis](#), [prod](#), [install](#), [test](#) jobs

Quick search

Job

Dataset

Task request

Task status

File

Summaries

Blocks: days

Errors: days

Nodes: days

[Daily usage](#)

Tasks - [search](#)

[Generic Task Reg](#)

[EvGen Task Reg](#)

[CTBsim Task Reg](#)

[Task list](#)

[New Tag](#)

[Bug Report](#)

Datasets - [search](#)

[Dataset browser](#)

[Aborted MC datasets](#)

[Panda subscriptions](#)

Datasets Distribution

[DDM Reg](#)

[Reg list](#)


Panda job error summary for last 24 hours (1.0 days)

All CEs and jobs. Show [production](#), [analysis](#), [test](#), all jobs/CEs

Job wall time: 317553 hrs Error losses: trans: 9971 (3.1%) panda: 8458 (2.7%) ddm: 3329 (1.0%) other: 1317 (0.4%)

Error type (type count)	Count	CPU-hrs	Latest	Code: Description
All	defined:708 assigned:264 waiting:0 activated:19020 sent:0 running:10202 holding:1599 transferring:5359 finished:40421			
	failed:8358 (17.1%)			
brokerageErrorCode (120)	120	0.0	09-17 13:11	100 : Unknown error code
ddmErrorCode (6)	1	0.0	09-16 18:14	100 : DQ2 server error
ddmErrorCode (6)	5	14.0	09-17 13:54	200 : Could not add output files to dataset
exeErrorCode (1114)	2	2.6	09-16 13:45	1101 : LRC registration error: Connection refused
exeErrorCode (1114)	1	0.9	09-16 20:13	1114 : Put error: Failed to import LFC python module
exeErrorCode (1114)	4	30.3	09-16 18:57	1131 : Put function can not be called for staging out
exeErrorCode (1114)	31	13.6	09-17 04:50	1132 : LRC registration error (consult log file)
exeErrorCode (1114)	7	14.8	09-17 10:25	1133 : Put error: Fetching default storage URL failed
exeErrorCode (1114)	1	26.2	09-15 10:22	1135 : Could not get file size in job workdir
exeErrorCode (1114)	875	7494.9	09-16 22:32	1137 : Put error: Error in copying the file from job workdir to localSE
exeErrorCode (1114)	13	159.2	09-16 15:34	1154 : Failed to register log file
exeErrorCode (1114)	6	58.6	09-16 15:20	1155 : Failed to move output files for lost job
exeErrorCode (1114)	1	11.8	09-14 15:01	1176 : Pilot has no child process
exeErrorCode (1114)	1	22.1	09-15 07:10	1211 : Missing installation
exeErrorCode (1114)	3	51.7	09-17 13:52	60000 : segmentation violation
exeErrorCode (1114)	117	399.1	09-17 13:44	60010 : segmentation fault
exeErrorCode (1114)	5	92.2	09-17 10:47	61200 : ServiceManager Unavailable
exeErrorCode (1114)	6	107.3	09-17 13:36	62600 : AthenaCrash
exeErrorCode (1114)	30	94.8	09-17 10:27	64100 : Transform output file error
exeErrorCode (1114)	11	52.2	09-17 12:15	69999 : Unknown Transform error

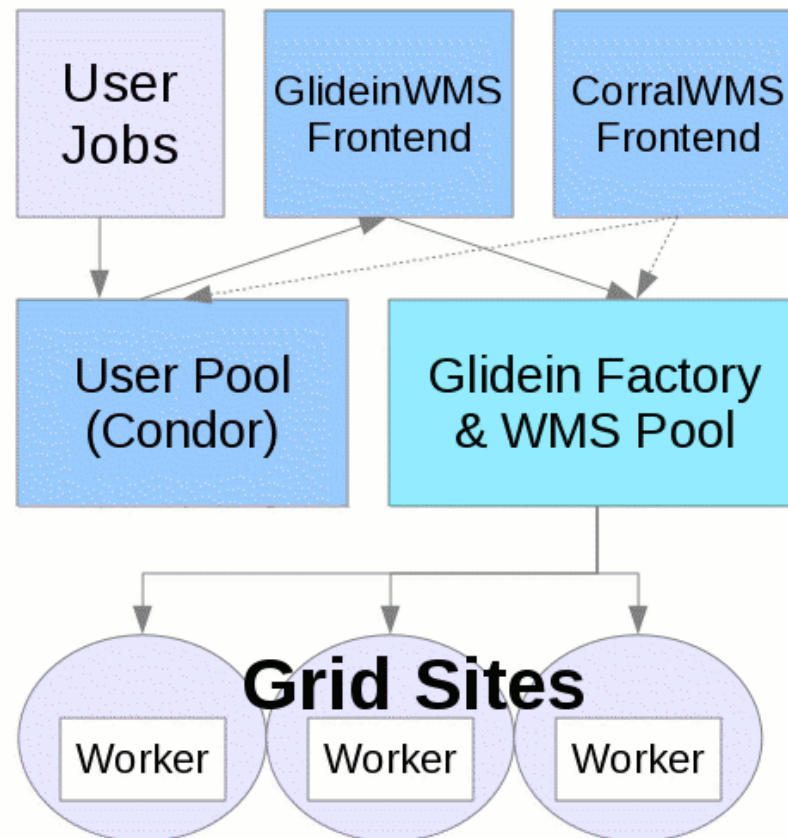
OSG Errors for period 09-08-2007



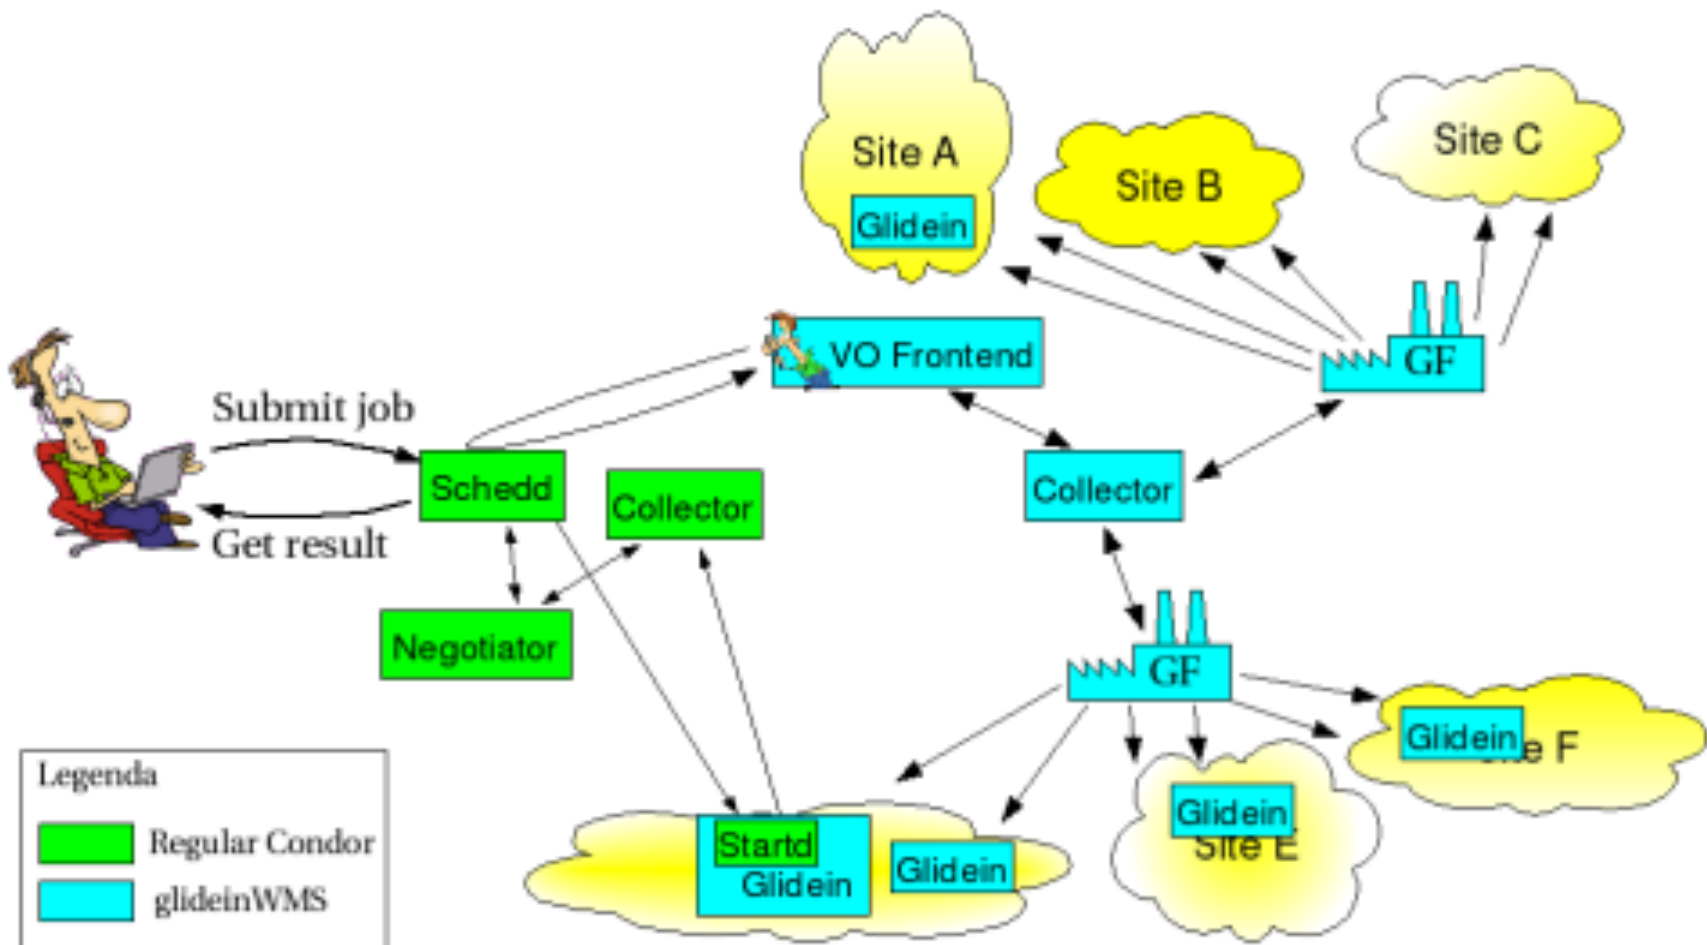
Glide-in WMS

- Provide a simple way to access the Grid resources. GlideinWMS is a Glidein Based WMS (Workload Management System) that works on top of Condor
- Workflow (as visible in the next page illustration)
 - Users submit jobs to the User Pool Condor schedd process.
 - The GlideinWMS Frontend polls the user pool to make sure that there are enough glideins (worker nodes) to satisfy user jobs. It submits requests to the glidein factory to submit glideins.
 - Alternatively, users can control their workflow with the CorralWMS frontend.
 - The glidein factory and WMS pool receives requests from the frontend(s) and submits a condor startd wrapper to entry points (grid sites).
 - The grid sites receive the jobs and starts a condor startd that joins the user pool. This shows up as a resource in the user pool.
- Final users can submit regular Condor jobs to the local queue

Glide-in WMS workflow



Glide-in WMS system



Submit File for Condor (or Campus Grid)

```
Universe = vanilla
#notify_user = <user email address>
Executable = serial64ff
#transfer_executable = false
# Files (and directory) in the submit host: log, stdout, stderr, stdin, other files
InitialDir = run_1
Log = serial64ff.log
Output = serial64ff.stdout
Error = serial64ff.stderr
#input = serial64.in
fetch_files=serial64ff.in
#Arguments = <arg1> <arg2> <argn>
should_transfer_files = IF_NEEDED
when_to_transfer_output = ON_EXIT
```

Queue



Submit File for Glidein WMS

```
Universe = vanilla
#notify_user = <user email address>
Executable = serial64ff
#transfer_executable = false
# Files (and directory) in the submit host: log, stdout, stderr, stdin, other files
InitialDir = run_1
Log = serial64ff.log
Output = serial64ff.stdout
Error = serial64ff.stderr
#input = serial64.in
fetch_files=serial64ff.in
#Arguments = <arg1> <arg2> <argn>
should_transfer_files = IF_NEEDED
when_to_transfer_output = ON_EXIT

# This line is required to specify that this job should run on GlideinWMS system
requirements = IS_GLIDEIN == True
x509userproxy=/tmp/x509up_u20003
```

Queue



?

!